Forecast-based Advance Release Flood Operations at Folsom Reservoir

Beth Faber, PhD, PE[1]; Marchia Bond[2]

[1] Hydraulic Engineer, Hydrologic Engineering Center, 609 Second St., Davis, CA 95616, Beth.A.Faber@usace.army.mil
[2] Hydraulic Engineer, Water Management Section, Sacramento District, 1325 J Street Sacramento, CA 95814-2922, Marchia.V.Bond@usace.army.mil

Abstract

Reservoir flood protection can be enhanced, without altering reservoir size or allocation of space, by freeing additional flood storage volume on an individual-event basis. Operators can evacuate water from the reservoir in advance of a flood by responding to National Weather Service (NWS) streamflow forecasts that warn of an upcoming event. This timely forecast information helps to increase the effective storage volume reserved for all purposes, including hydropower generation and water supply for agricultural, municipal and environmental users.

Folsom Reservoir on the American River benefits from sophisticated short-term forecasting from NWS that provides not only a 5-day streamflow forecast, but an estimate of forecast uncertainty. A strategy for implementing a forecast-based Advance Release is being considered to let operators make efficient use of the forecast information to provide additional flood protection. This forecast-based release can also create additional conservation storage due to the ability to evacuate that space in advance of a flood event.

Advance Release strategies considered at Folsom are focused on avoiding impact to other users of the reservoir and river. Negative impacts studied are those resulting from a “false alarm” forecast that may lead to loss of reservoir storage or higher than necessary reservoir release. Strategies involve achieving additional flood storage while maintaining a minimal probability of negative impact.

Introduction

The flood protection afforded by a reservoir can be enhanced by providing (or freeing) additional reservoir volume for storage of flood water. Additional volume can be made available on an individual-event basis by responding to NWS streamflow forecasts that warn of an upcoming event. Heeding this forecast information, operators can initiate a preemptive release to evacuate water from the reservoir in advance of the flood. Figure 1 depicts a reservoir routing a flood event, first with normal flood operations (a), and then with Advance Release (b). The preemptive release makes the later high release unnecessary.

This idea is discussed and demonstrated for Folsom Reservoir on the American River in California. An Advance Release strategy was developed to provide additional flood protection by making efficient use of both the additional release capacity soon to be provided by enlarged outlets and the American River streamflow forecast information generated by the California/Nevada River Forecast Center (CNFRC).

Folsom Reservoir has a flood pool large enough to manage most flood events. Advance Release is intended for those occasions in which the event is larger than can be managed by the current flood pool/enlarged outlet combination. The flood events that would seriously challenge the reservoir with its normal operating rules are larger than any that have yet occurred on the American River within the historical streamflow record. Such events would have somewhat less than a 1%-chance of occurring each year.

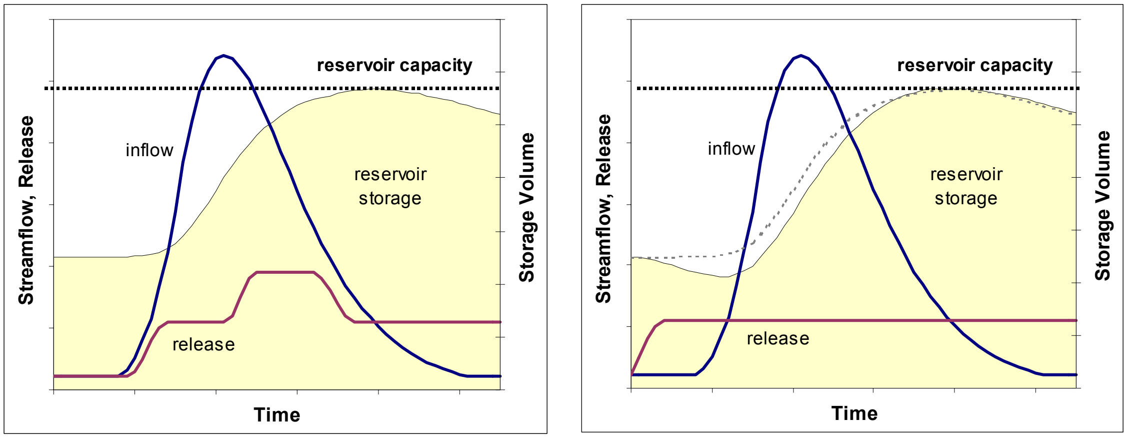

(a) without Advance Release (b) with Advance Release

Figure 1. Operation of a reservoir, without and with Advance Release

Preemptive release of water in preparation for a flood event is not a standard practice at facilities operated by the USACE. To avoid increasing flood protection at the expense of other functions of Folsom Reservoir, advance releases must not adversely impact those functions. An Advance Release activity could not be considered successful unless it avoided the following outcomes:

1. Release of a higher flow *in advance* of the event than would have been released *during* the event with no preemptive action.
2. Failure to refill the reservoir's conservation pool by the end of the event, causing Advance Release to interfere with conservation storage and hydropower generation.

Either outcome would suggest that the Advance Release had been unwarranted or ill-advised in that instance. However, when decisions are made based on uncertain information such as streamflow forecasts, some risk of negative outcome is always incurred. The benefit gained from such decisions often outweighs the risk. This study investigated the Benefit/Risk relationship for Advance Release at Folsom Reservoir, based on the impacts listed above. Other impacts, such as a decrease in the reservoir's cold water pool, were discussed but not explored in this phase.

Operating Reservoirs with Uncertain Streamflow

Currently, decisions about Folsom Reservoir operation are based on current reservoir inflow and future inflow derived from *gauged* upstream flow and precipitation, providing 4 to 12 hours lead-time. This paper addresses the proposal to make operating decisions based on 5-day streamflow forecasts built with *forecasted* precipitation. Such forecasted streamflows are much less certain, incorporating the uncertainty of precipitation forecasts as well as of hydrologic modeling.

Good decision-making requires information about factors outside of our control that nonetheless have impact on the outcome of our decisions. Decisions based on uncertain information such as forecasts require an understanding of the probabilistic nature of the information.

Consider the operation of reservoirs that provide both flood protection and water supply storage. These two functions are by definition in conflict, each requiring reservoir volume but using it in the opposite way. For water supply, volume is used to store incoming water for later use, while for flood protection that volume must remain empty and available to contain flood flows when they arrive. Serving both of these functions requires a tradeoff between them that is defined by the target storage level of the reservoir (the guide curve).

In reservoir operation, the main factor outside our control is the streamflow into the reservoir. Given *no* information about future streamflows (ie, complete uncertainty), it would be necessary to maintain separate reservoir volumes for each function, with the water supply volume allowed to remain as full as possible, and the flood protection volume kept empty between flood events. The benefits attainable by each function are proportional to the volume allocated to it.

If, on the other hand, we had perfect and unlimited foresight of streamflows (ie, no uncertainty), each reservoir function could make use of the total volume of the reservoir. Water stored for supply could be evacuated in time for a flood event—the entire reservoir if necessary—and be completely refilled by the event itself. With the total reservoir volume cooperatively serving both functions, the benefits attainable by each function are proportionally greater.

Actual reservoir operation is somewhere between these two extremes. Any information used in decision-making decreases uncertainty and brings us closer to the perfect-foresight ideal. However, using uncertain information introduces some risk of making the "wrong" decision, a decision with an outcome worse than one based on no information at all.

Use of Information. Any hydrologic information provides increased certainty about the magnitude and the arrival of flood events. For example, the historical record of streamflow at a site is the most straightforward information available to decision-makers. A frequency analysis of annual flood peaks and volumes addresses probable flood magnitude, allowing the flood pool to be sized for a given probability of being overwhelmed by a too-large event (ie, some risk of a too-large event remains.)

The historical record also contains information about *when* flood events occur. Considering the monthly occurrence of extreme precipitation events, we can define a likely flood season and allocate flood protection volume during only the part of the year when a flood event is probable. During the remainder of the year, that reservoir space can be used to store water for supply. There remains a risk of a flood occurring when it was improbable to do so, and therefore without flood storage space available.

The availability of more current hydrologic information further decreases uncertainty and allows more precise operations. When timely information (such as climate indicators or soil moisture readings) is available, flood reservation can be adjusted within-year based on the updated conditional probabilities of flooding. This adjustment allows only the necessary reservoir volume to remain empty for flood protection, but again with some risk of an unexpected outcome.

As the time horizon for decision-making decreases, the uncertainty in the size and arrival of the next flood event decreases. Near-term weather forecasting provides more accurate information about imminent streamflow, and rather than allowing for the large variance of annual flood probabilities, or even conditional probabilities specific to the current year, operation can be tailored to a more specific estimate of the approaching storm. This information allows the flood pool to be enlarged at need, and might allow additional conservation storage knowing that space could be evacuated before an event.

Table 1 summarizes the information available at long, middle, and short time horizons, and how this information is used to aid reservoir operation.

Table 1. Use of Information in Long-, Mid- and Short-term flood operations

	Long-term *year to year*	**Mid-term** *current year*	**Short-term** *current week*
Available Information	**Historical record** of precipitation and extreme flows (peak and volume)	Annual **climate indicators** (El Ni⋆o, etc), basin wetness, upstream storage	**5-day**, 6-hourly streamflow **forecasts**
Processing of information	• Development of peak and volume **frequency curves** to determine the probability of events of any magnitude • **Definition of flood season**, based on the seasonal flood risk throughout the year.	• Determination of the appropriate level of ***variable* flood space**, based on likely flood magnitude, basin wetness, upstream storage levels, etc.	• Real-time use of the best-estimate streamflow forecast **to simulate and adjust operations** • Development of probability **distribution of** imminent **event volume** to compute Advance Release
Use of information in Reservoir Operation	• **Size the reservoir flood pool** to contain an event of the selected frequency. • **Maintain the empty flood pool** when flood events are highly probable.	• **Maintain the variable flood pool** of the size determined above.	• Make more **decisive releases** as the event nears and forecast uncertainty decreases.

Ways to be "Wrong." The most efficient decisions in the face of uncertainty are made with an understanding of that decision's consequences given any possible outcome. In evaluating Advance Release as a feasible operating strategy, we must be aware of how decisions can be "wrong" (ie, which outcomes would lead to a negative result) and what the consequences of those outcomes would be.

In the simplest sense, forecasts can be off-target by either overestimating or underestimating an imminent flood event. If forecasts underestimated a large event, no Advance Release would be made and an opportunity to reduce flood damage would be missed. The opposite error would be forecasts that overestimated a smaller event, prompting an unneeded Advance Release. Release based on an overestimate might be larger than would have been necessary had operators responded to the event without the use of flood forecasts, and perhaps too large to recover the reservoir drawdown, impacting water supply, hydropower, and environmental uses. Due to the injury to water users who do not necessarily share in the flood damage reduction benefits of Advance Release, this type of error is an important concern in considering Advance Release.

These ideas lead to a matrix of forecasts, actions, outcomes and consequences. Floods are divided into three levels to capture the situations relevant to Advance Release, defined by whether the event needs, and can recover from, Advance Release. ***Very Large*** events would require Advance Release to avoid release above channel capacity. ***Large*** events do not require Advance Release, but the reservoir would refill if Advance Release were initiated. ***Small*** events would be unable to refill after Advance Release.

Table 2 details the consequences of forecast-based decisions, based on the above definitions. The intersection of column (forecast) and row (occurrence) represents the pairing of that forecast with that subsequent occurrence. Because only *very large* forecasts suggest Advance Release, that column is shaded to highlight an action taken.

Table 2. Possible Combinations of Flood Forecast and Flood Occurrence

	Forecast of Very Large Event	**Forecast of Large Event**	**Forecast of Small Event**
	Forecasts Suggest Advance Release	***Forecasts Do Not Suggest Advance Release***	
Very Large Event Occurs	*Reduced Flood Damage*	Flood Damage (missed opportunity)	
Large Event Occurs	No Flood Damage or Impact	*Not an Advance Release situation*	
Small Event Occurs	Impact on Supply		

The upper left and lower right corners of the grid are situations in which the forecasts were correct in at least the general size of the event. The "*very large* forecast / *very large* occurrence" combination is the one that drives the Advance Release strategies explored here, with a focus on minimizing the impact of mistakenly using Advance Release when it is not needed ("*very large* forecast / *large* or *small* occurrence").

These ways to be "wrong" have been defined with a yes/no, on/off description of Advance Release, based on a fairly simplistic all-or-nothing definition. Because a goal of effective decision-making under uncertainty is to minimize the negative impact of being wrong, the Advance Release strategies in this paper are designed with an effort at robustness that is *not* all-or-nothing, but instead tailored to forecasts and responsive to forecast updates. The situation is therefore not as straightforward as these figures imply.

Advance Release Strategy Components

There are several components to an Advance Release strategy, each of which influence the strategy's overall effectiveness. Options for each component are outlined here for discussion, although at this stage no recommendations are made. The components are:

(A) **lead-time** at which Advance Release action may be taken,
(B) **trigger for initiating** Advance Release procedures,
(C) **release value implemented**, and how it is determined, and
(D) **trigger for discontinuing** Advance Release when the event forecast decreases

Lead-Time for Action. The National Weather Service (NWS) produces 6-hour time-step streamflow forecasts for the American River that extend 5 days into the future.

Logically, uncertainty in the forecast is less for shorter lead-time than longer, ie, less for the forecast of tomorrow than of 3 days from now. Representatives of the NWS CNRFC have advised that decisions based upon forecasted flows more than 3 days distant would be ill-advised, but incorporating forecasted flows 3 days distant or less would be reasonable. Therefore, these Advance Release strategies assume action will be taken based upon forecasted streamflow levels 3 or fewer days in the future.

Triggers to Initiate Advance Release. The trigger to initiate Advance Release is the magnitude of forecast for which we feel action must be taken. This trigger represents the decision-makers' level of risk-aversion against failing to initiate Advance Release when needed, or initiating when not needed. A smaller trigger reduces the likelihood of missing an opportunity to reduce flood damage with Advance Release, but would increase the likelihood of acting on a false alarm, while a more extreme trigger would provide the opposite result.

Advance Release triggers discussed here are based on the 5-day streamflow forecast for the American River provided by the NWS CNRFC. There are several possible elements of the forecast that can be used to trigger action:

1) *A forecasted peak streamflow of a certain value*, perhaps 300,000 cfs.
2) *A forecasted event volume of a certain value*, perhaps 1,000,000 acre-feet.

These two triggers might be applied together with an "OR" condition. In other words, if the forecast predicts *either* a peak flow greater than 300,000 cfs *or* an event volume greater than 1,000,000 acre-feet, then a release would be initiated.

3) *A particular quantile of the forecasted event satisfying either trigger 1 or 2*. For example, the 25% chance exceedance (75% quantile) forecast hydrograph meets one of the above triggers.
4) *The forecasted event (or a particular quantile of the forecast) causing a routed reservoir release greater than channel capacity, without Advance Release*.
5) *A particular probability of release exceeding channel capacity without Advance Release*, perhaps 15%. Two methods of computing this trigger are:
    a. Route several quantiles of the hydrograph though the reservoir, and note the smallest to produce a release over channel capacity. If the 85% quantile (15% chance of exceedance) is the smallest to cause release greater than channel capacity, probability of exceeding channel capacity would be 15%.
    b. A perhaps more rigorous method is to use the "uncertainty version" of the Folsom Reservoir Release Forecast Model (RRFM) developed by Utah State University. RRFM develops thousands of likely realizations (traces) of the forecasted event that are each routed through the reservoir to produce release traces (Bowles, 2001). The probability of exceeding channel capacity is computed simply as the relative frequency of release traces that exceeded channel capacity.

Each of these triggers have merit, some for simplicity, some for capturing essential aspects of the flood-management problem.

Release Implemented. The Advance Release can be more or less dependent on the forecast, based on the desired sensitivity of the strategy to forecast magnitude and error. A strategy less dependent on the forecast would be less sensitive to forecast error, but would also suggest a release that was less appropriate for the imminent flood event. A strategy more dependent on the forecast would make better use of that information, but would likewise be more sensitive to error.

Increment-based strategy — Less Dependence on Forecast. An Advance Release strategy with less dependence on the forecast is a release that is some increment above inflow (allowing reservoir storage to decrease), with a maximum release of channel capacity. For example, given increment of 25,000 cfs, an inflow of 35,000 cfs would suggest a release of 60,000 cfs.

Once Advance Release has been triggered, and for the duration of time it remains triggered, the release recommended by this strategy would not depend upon the actual streamflow forecast, and so would not be sensitive to errors in the forecast. The relevant (and adjustable) parameter of this strategy is the increment of release above inflow. A similar strategy might begin with a smaller increment, and increase it as the event grew nearer, considering that forecast uncertainty decreases with successive forecasts.

Volume-based strategy — More Dependence on Forecast. A strategy with more dependence on the flood forecast is one designed to manage the volume of the forecasted hydrograph. The computed release is based on the volume of water that "will not fit" in the reservoir flood pool. As successive event forecasts varied in magnitude and volume, the computed release would also vary, making this strategy sensitive to forecast error.

Folsom Reservoir has a reserved flood pool able to store a certain volume of flood water, $V_{floodpool}$. Considering flood operations in the simplest light, an event producing a volume smaller than $V_{floodpool}$ (or the currently empty space in the reservoir) would not require *any* release during the event, while an event producing a volume larger than $V_{floodpool}$ *would* require release. The minimum volume of release necessary to manage the event would be the inflow volume in excess of $V_{floodpool}$, depicted in Figure 2. Completely efficient flood operation would limit total release to that minimum volume, maximizing use of the flood pool for flood event storage and so minimizing peak release. In such completely efficient operation, the flood pool fills just as inflow falls below safe release.

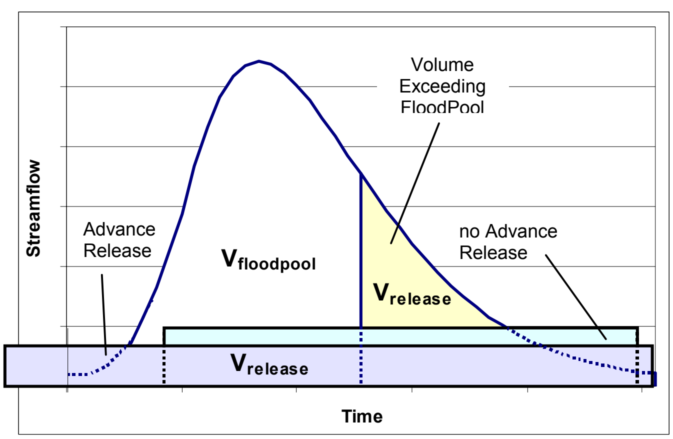

Figure 2. Minimum Release during a Flood Event, Level of release vs duration of release

These ideas have addressed the *volume* of water released during a flood event. The release *rate* is equal to that volume divided by the duration of release. Advance Release lengthens the period of time (by starting earlier) water is released from the reservoir to pass the flood event. Given the required volume, the release rate will be smaller over a longer period of time, as depicted in Figure 2. Note that starting release earlier means release will initially be higher than inflow, lowering the reservoir's storage volume *before* the event. Increasing the event volume that will "fit" in the reservoir serves to decrease the required release *during* the event.

Based on these ideas, a strategy for volume-based Advance Release considers the volume of the latest inflow forecast and the storage space currently available in the reservoir, and releases the difference between these volumes. A release rate would be computed for each updated inflow forecast as follows:

(1) From the latest streamflow forecast, compute both the volume of inflow expected during the flood event (**EventVol**) and the length of the event (**EventTime**) (from forecast date/time until inflow falls below a safe release rate.)
(2) Compute the amount of flood storage space available in the reservoir. (**SpaceAvail**)
(3) The forecasted inflow volume (**EventVol**) minus the available storage space (**SpaceAvail**) is the minimum volume to be released during the event (**ReqReleaseVol**).
(4) Required release is computed as **Release** = **ReqReleaseVol** / **EventTime**.

It is clear that the earlier the release is computed and initiated, the more time is available for release (ie, **EventTime** is larger) and the lower the computed release will be. Therefore, action taken farther in advance of the event peak may be less extreme (a lower release) than action that waits for a shorter lead-time.

Maintaining "No Impact" on other Folsom Objectives (the Refill Constraint). One of the initial stipulations for an Advance Release strategy at Folsom Dam was that it not impact the reservoir's other uses. Impact could be caused by heeding a "false alarm" forecast that greatly overestimated a flood. The reservoir drawdown in preparation for the forecasted flood then might not be refilled, leaving the conservation pool depleted and so affecting water supply and hydropower. ("Impact" is therefore defined here as the reservoir ending the event below the Guide Curve.) It is difficult to ensure no impact when considering action based on uncertain information because there is always some probability of impact if an Advance Release is made. Instead, the probability of impact can be limited to some reasonable value.

For the reservoir to at least recover the Advance Release drawdown, total release volume must be less than or equal to total inflow volume. If the projected volume of release is less than a volume that has a 99% probability of being exceeded as flood event inflow (Figure 3, derived from forecast standard error), there is a 99% probability that the reservoir will at least refill after Advance Release, or a 1% probability of impact in the case of a "false alarm."

The aggressiveness (and perhaps effectiveness) of the Advance Release strategy varies inversely with the probability of reservoir refill it provides. In other words, a larger Advance Release would achieve greater flood protection, but would also run a higher risk of not refilling after a "false alarm" forecast. In this study a 1% probability of impact (99% refill probability) was chosen in the Advance Release strategy to satisfy "no impact" as closely as is reasonable. Sensitivity to that choice was examined in simulation by varying this target refill probability.

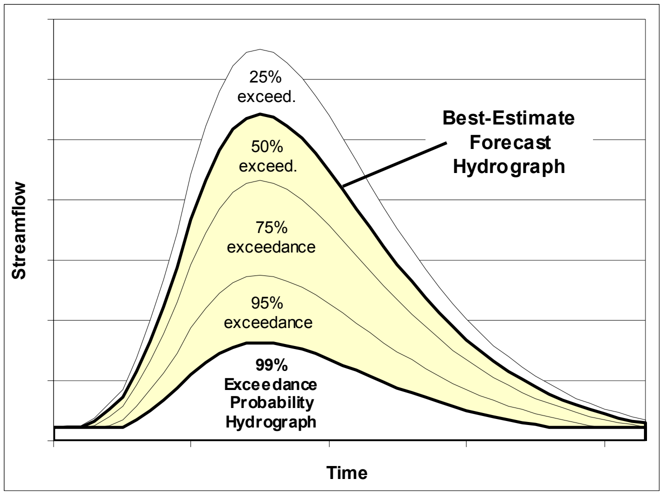

Figure 3. Best-Estimate Inflow Forecast and Quantiles, based on Forecast Uncertainty

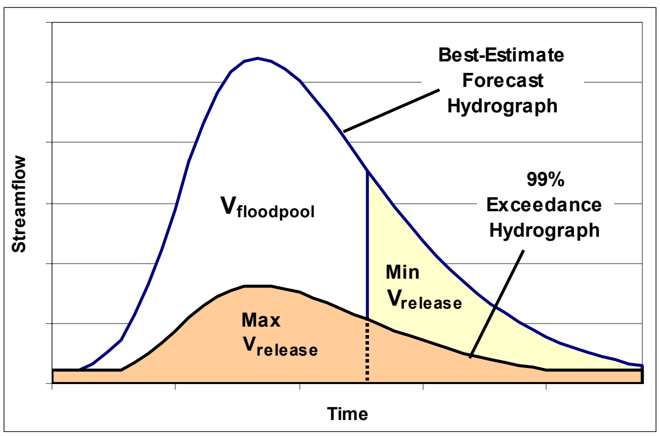

Figure 4. Comparison of Minimum Release Volume to Maximum Release Volume

The volume-based strategy considers the volume of release required to manage a flood event. With the best-estimate forecast of event inflow, the strategy computes the portion of inflow volume that exceeds available flood storage space to define the ***minimum*** that must be released to avoid overtopping the dam. With the incorporation of the refill constraint (release volume ≤ inflow volume), the ***maximum*** volume that may be released to limit risk of impact to 1% is the 99%-exceedance inflow volume, ensuring release can be replaced with at least a 99% probability. Figure 4 displays the minimum required release volume and the maximum allowable release volume.
As long as the maximum allowable release volume is larger than the minimum required release volume, the refill probability constraint does not effect the release. However, when forecast uncertainty is large (and so the 99%-exceedance hydrograph is small), the maximum allowable release may be smaller than the minimum required. The question of how to proceed is an important one. In this study, the maximum release volume to avoid impact was maintained, rather than the minimum required volume. This choice was made because Advance Release is an *added* element of flood protection that will be followed by standard flood operations when the flood event arrives.

Trigger to Discontinue Advance Release. A robust Advance Release strategy that minimizes the consequence of relying on an incorrect forecast must be able to adjust (or retreat from) a “wrong” decision. It must be possible to un-trigger or reduce the Advance Release if necessary. Streamflow forecasts tend to vary above and below the actual occurrence, however. Were Advance Release discontinued when the latest streamflow forecast predicted a less severe event, Advance Release could start and stop as forecasts moved above and below the trigger. But as forecasts decrease, at some point the Advance Release must respond.

For the increment-based Advance Release strategy, in which the release value is not directly dependent on the forecast, the "stop-trigger" can perhaps be similar to the "start-trigger," but lower. If a forecasted peak of 300,000 cfs *or* volume of 1,000,000 acre-feet is required to initiate Advance Release, then a peak less than 200,000 cfs *and* volume less than 600,000 acre-feet might trigger halting the Advance Release. The values should be chosen carefully, however this study did not focus on the increment-based Advance Release strategy, and so these values have not been evaluated.

For the volume-based strategy in which computed release is dependent upon the volume of the forecast, a "stop-trigger" is less necessary. The computed release value is tailored to each new forecast, and directed at evacuating the necessary volume of water. The Advance Release does not need to be *discontinued* to avoid evacuating too much water, as the computation responds to the forecast and will suggest a smaller release.

Simulation and Results

The Forecast/Occurrence combinations simulated in this study are those that activate a forecast-based strategy. These are the situations in which forecasts of a "*very large*" event (which would require Advance Release to be managed successfully) suggest implementing Advance Release, and include the cases in which an event of that size does not occur, as shown in the shaded column of Table 3.

Because the performance of a forecast-based Advance Release strategy depends so heavily upon the series of forecasts leading up to a given event, simulating reservoir operation for that event with only one set of forecasts would not provide a complete assessment. To determine a realistic measure of effectiveness, the strategy must be simulated with many realizations of the forecast time-series to evaluate the range of possible outcome of that event.

In this study, the "*very large* forecast / *very large* occurrence" combination was repeatedly simulated with 120 possible realizations of the forecast time-series, which were generated artificially. This exercise allowed probabilistic evaluation of the Advance Release strategies when Advance Release was truly needed. However the difficulty with generating "false alarm" forecasts made it difficult to evaluate the likelihood and effect of Advance Release when it is not only not truly needed, but in fact harmful, i.e., "*very large* forecast / *small* occurrence." As this is the combination that would cause impact to other users of Folsom Reservoir, this lack of analysis means that this preliminary study is not yet completely thorough. The NWS CNRFC is currently soliciting a study to reproduce the historical forecast series between 1965 an 2000 with current forecast methods and models. This effort will perhaps provide the likelihood information to more completely evaluate this impact.

Events Simulated. To produce an Advance Release strategy capable of responding successfully to a wide range of hydrology, it is necessary to work with a group of events as varied as those the reservoir is likely to experience. The events simulated in this study were scaled-up versions of the 1986 and 1997 flood events, and the American River Project $^{1}/_{150}$-exceedance (150 year) Design Storm. The 1986 event produced a large volume over several days without having an extremely high peak, while the 1997 event was shorter in duration with a higher peak. The American River Project design storm includes a small initial flood wave, allowing analysis of consecutive flood events. These different types of storms present different operating challenges to the reservoir. The historical events were scaled-up to force the reservoir to release above channel capacity with normal flood operation (150% for the 1997 event, and 170% for the 1986 event). Figure 5 displays the event hydrographs simulated in this study and Table 3 contains summary values.

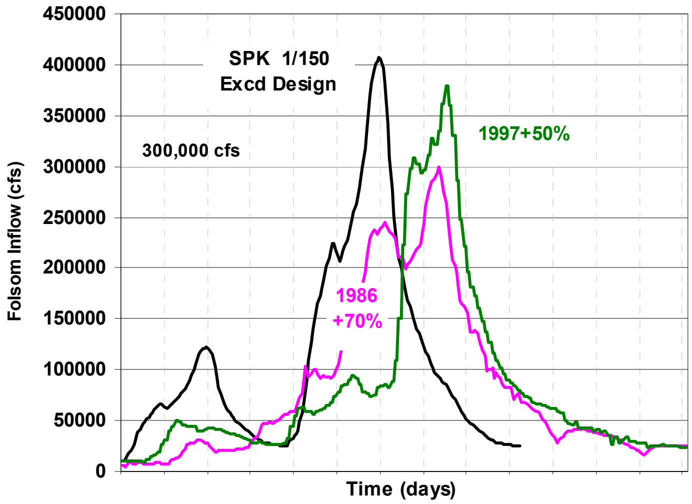

Figure 5. Three Events Simulated

Table 3. Event Volume and Peak Flow for Flood Events

	Total Event Volume (af)	Peak Flow (cfs)
Actual 1997 Event	1,135,517	252,538
1997 Event + 50%	1,703,275	378,807
Actual 1986 Event	998,257	176,080
1986 Event + 70%	1,697,037	299,336
$^{1}/_{150}$% Exceedance Design Storm	1,911,386	406,640

Strategies Simulated, Triggers Used, Performance Indicators. Both strategies introduced were simulated. The increment-based strategy was evaluated with increments of 25,000 cfs and 50,000 cfs above inflow, and the volume-based strategy was evaluated with target refill probability levels of 99%, 95%, 90%, 80%, 70%, 60%, and 50% to determine sensitivity to those choices. Of the various Advance Release triggers discussed, a combination of (1) and (2) was used in the simulations, as was trigger (4). Results showed little difference between these triggers, not visible on graphical plots.

The two strategies were evaluated based on average reservoir drawdown achieved, the average maximum release and the probability of release exceeding the 115,000 cfs channel capacity. (Averages were taken across the 120 simulations of each event, based on different sets of artificial forecasts.) These factors address the strategy effectiveness when the event is of a size to make Advance Release necessary. A further factor for comparison is how the strategies behave when forecasts over-estimate the oncoming event, threatening impact on water supply and power generation. To consider the strategies' performance when Advance Release is not actually required, we note the reservoir refill probability maintained by the releases. With the volume-based strategy, the minimum refill probability is maintained directly. With the increment-based strategy, the refill probability based on the forecast uncertainty (averaged over all 6-hour forecasts) is computed as an additional step.

For comparison of various strategies, a level of 500,000 af was chosen for the variable flood pool. The initial reservoir condition was the bottom of the flood pool. Both Advance Release strategies were simulated for each event using each set of forecasts, each level of refill probability, both triggers. (This array required 2 strategies * 3 events * 120 forecast sets * 7 refill probabilities * 2 triggers = 10,080 simulation runs.)

Results for the 1997 event are included in this paper, and the remaining results can be found in (USACE, 2002).

1997+50% Event, Volume-based Strategy. Figures 6 displays an example of the use of Advance Release with a single set of forecasts, showing the 1997 flood event scaled up by 50% and the reservoir operating with normal flood operations, and with Advance Release operations.

The most significant difference between the reservoir routings without and with Advance Release is that the *without*-routing is directed by the ESRD to release a maximum of 155,000 cfs, while the *with*-routing keeps the maximum release much closer to safe carrying capacity of 115,000 cfs. With Advance Release, reservoir storage was decreased by 61,000 af during the Advance period, providing more storage space to manage the flood event. This lower storage level served to prevent a storage/inflow combination that causes the ESRD to direct an extremely large release. A similar routing maintaining a less restrictive 95% refill probability avoided activation of the ESRD and release greater than 115,000 cfs.

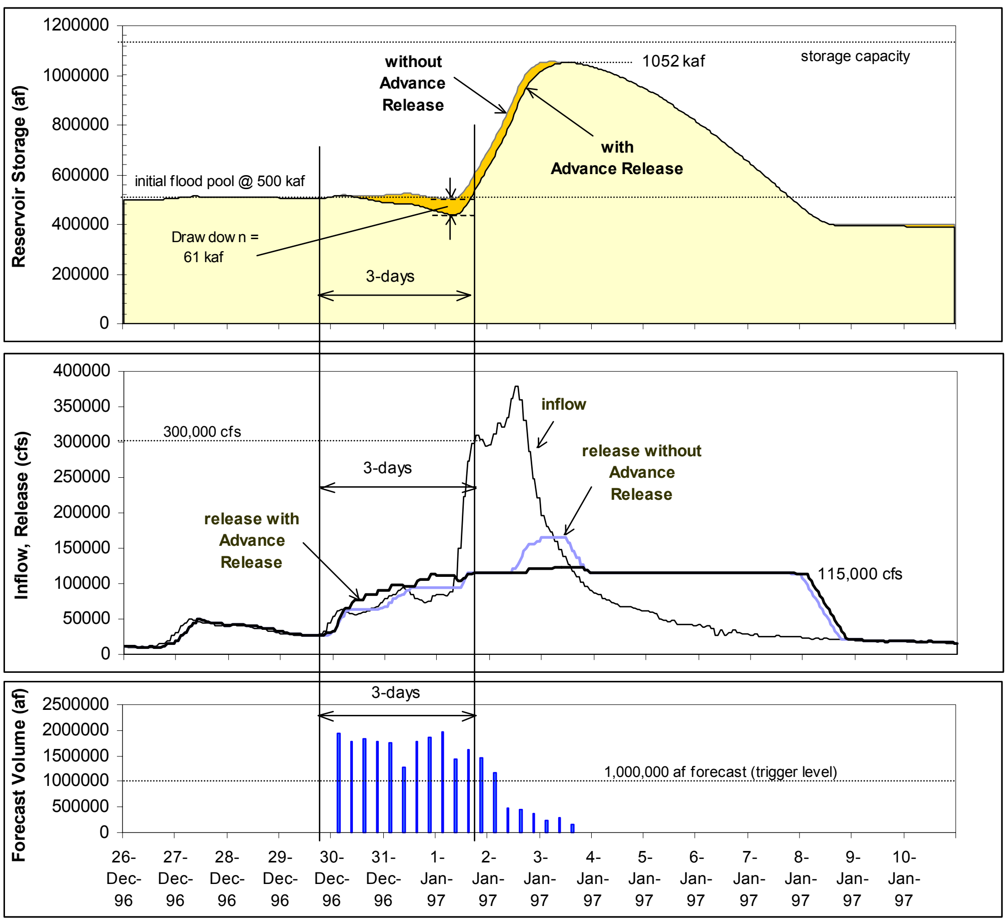

Figure 6. 1997 + 150% Event, without and with Advance Release

Although illustrative, the outcome of Advance Release based on one set of forecasts was not sufficient for evaluation of the event or the strategy. Further simulations routed this event using 120 sets of artificial forecasts, and are summarized by reporting average volume of drawdown and the frequency of release exceeding 115,000 cfs during the event. Figure 7 displays these results, showing the sensitivity of strategy effectiveness to the required refill probability. The figures also include computed refill probabilities maintained by routings using the increment-based strategy (circular markers). The strategies perform similarly when the flood event occurs as expected.

Figure 7 shows that 99% refill probability is restrictive enough to limit the effectiveness of Advance Release as compared to the lower values of refill probability. It also shows that reducing refill probability to 90% achieves nearly the entire effect of the lower values. From this comparison, it seems that consideration of refill probability values can be limited to the range between 90% and 99%.

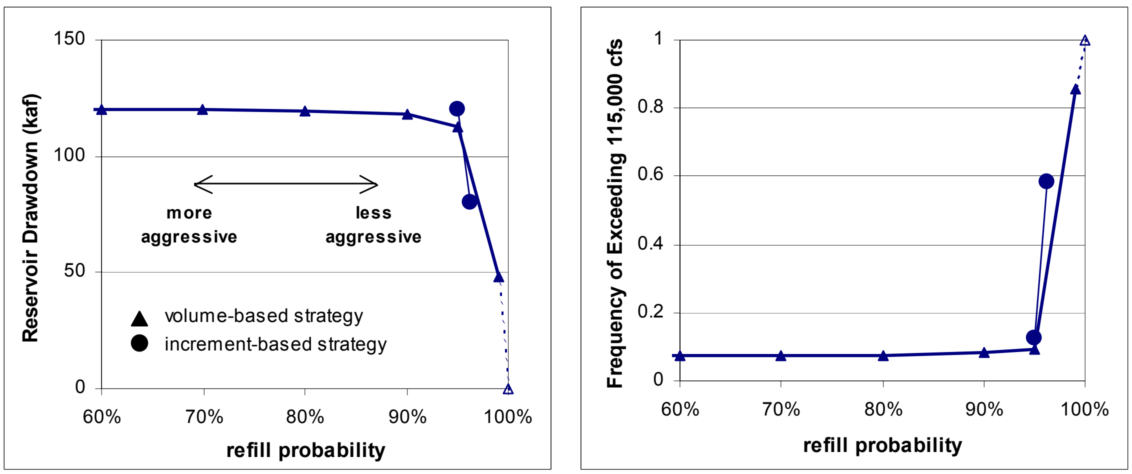

Figure 7. Drawdown and refill probabilities for both strategies, 1997+50% event

Performance of Volume-based Strategy with "False Alarm" Forecasts. Simulation performed has been for cases in which an event needing Advance Release *has* occurred. The more aggressive strategies therefore look the most promising. We are also concerned with forecasts predicting an event requiring Advance Release, and then a much smaller event occurring. This is the case that might lead to failure to refill drawdown, impacting water supply and hydropower generation. Simulation of this situation is of interest, but generation of artificial forecasts for this scenario is difficult.

The artificial forecasts generated for the earlier case of "the event occurs" were based on the conditional probability distribution for P[forecasts | event occurs] (*read as "probability of this forecast given that event occurs"*). This distribution seems fairly well defined, and the historical example of 1997 was available for analysis. False alarm forecasts, however, would be based on the conditional probability distribution for P[forecasts | event *does not* occur]. No historical data for this situation has been archived, although it is possible there have been examples. The relationship between the actual streamflow and potential large forecasts would be less well-defined, and many historical or synthetic examples would be needed to approximate it.

The NWS CNRFC and Hydrologic Research Center (the group responsible for developing the error propagation component of the NWS CNRFC American River forecast model) is recreating an historical series of streamflow forecasts between 1965 and 2000. Several realizations of this period-of-record series will be created, representing differing choices made by forecasters in using the array of climate and atmospheric models at their disposal. The resulting sample of forecasts, paired with the actual streamflows of that time period, will form an extremely useful dataset for evaluating the properties of streamflow forecast errors. This data will allow us to assess the probability of negative impacts from various Advance Release strategies, and test the effectiveness of the refill probability constraint in the volume-based strategy.

Without this information yet available, the analysis of the strategies' response to false alarms was postponed until the next phase of the study.

Conclusions and Recommendations

This study demonstrated the potential effectiveness of a forecast-based Advance Release, and described various strategies for determining the timing and magnitude of an Advance Release. We see that the potential effectiveness of these strategies must be balanced against the risk of impacts to other reservoir objectives in the case of substantially incorrect streamflow forecasts. While various Advance Release strategies and component options were presented and their relative effectiveness explored, no choice was made among the various options. Those decisions will be made by SPK Water Management, based upon the operational, institutional and political realities of undertaking a forecast-based operations scheme, and further study to more completely evaluate the parameters of chosen options.

Potential Operational Flexibility. As stated earlier, it is difficult to balance the benefits and the costs of Advance Release because often those benefits and costs accrue to different interests/stakeholders. The benefits are experienced by those who receive flood protection from Folsom Dam, and the possible costs would be felt by those using Folsom for other objectives such as water supply and hydropower generation. Because of this disparity, the benefits cannot be considered to outweigh the costs, even if they greatly exceed them.

A benefit/cost balance might be achieved if benefits can be afforded to those who bear the risk of Advance Release based on incorrect (overestimated) forecasts. One suggestion discussed by the multi-agency Advance Release working group was a policy to allow operational flexibility in the form of water storage in the flood pool during the winter flood season, with the understanding that the surplus storage would be evacuated after a flood forecast. Such operational flexibility could go a long way in gaining support for Advance Release and realizing the potential benefits of this approach.

References

Bowles, D., 2001, personal communication about Reservoir Release Forecast Model developed by Utah State University.

USACE, 2002, "Forecast-Based Advance Release at Folsom Dam: Effectiveness and Risks – Phase I," Hydrologic Engineering Center, PR-48.